

REMARKS

The above amendments to the above-captioned application along with the following remarks are being submitted as a full and complete response to the Official Action dated December 7, 2006, and the telephone interview with the Examiner and her supervisor on March 6, 2007. Applicants thank the Examiners for taking the time to conduct the telephone interview.

The Examiners indicated the term "suspension period" is not clearly defined with respect to the steps recited in the after-failure synchronizing procedure, and the Examiners suggested directly reciting the remote copy manager (RCM) 105-2 (which resides in the host computer 101-2) in conjunction with the steps in Fig. 6A. Such claim amendments probably will overcome the prior art of record, but these features should be clearly recited in the claims to be formally considered after the response is filed.

In view of the above amendments and the following remarks, the Examiner is respectfully requested to give due reconsideration to this application, to indicate the allowability of the claims, and to pass this case to issue.

Status of the Claims

Claims 1-21 are under consideration in this application. Claims 1, 8 and 15 are being amended, as set forth in the above marked-up presentation of the claim amendments, in order to more particularly define and distinctly claim Applicant's invention.

The claims are being amended to correct formal errors and/or to better recite or describe the features of the present invention as claimed. All the amendments to the claims are supported by the specification. Applicant hereby submits that no new matter is being introduced into the application through the submission of this response.

Formality Rejection

Claims 1-21 were still rejected under 35 U.S.C. §112, second paragraph, on the grounds the term "periodically" is indefinite in claims 1, 8 and 15.

As indicated, the claims are being amended to delete the term "periodically". Accordingly, the withdrawal of the outstanding informality rejection is in order, and is therefore respectfully solicited.

Prior Art Rejections

Claims 1-2, 5-9, 12-16 and 19-21 were rejected under 35 U.S.C. §103(a) as being unpatentable over US Patent No. 5,555,371 to Duyanovich et al. (hereinafter “Duyanovich”) over US Pat. No. 5,051,887 to Berger et al. (hereinafter “Berger”) and a newly cited reference US Patent No. 6,578,120 to Crockett et al. (hereinafter “Crockett”), and claims 3-4, 10-11 and 17-18 were rejected under 35 U.S.C. §103(a) as being unpatentable over Duyanovich in view of Berger, Crockett and US Patent No. 6,938,045 to Li et al. (hereinafter “Li”). These rejections have been carefully considered, but are most respectfully traversed.

In a system including a plurality of primary storage subsystems 102 #1, 102 #2, a plurality of secondary storage subsystems 102 #3, 102 #4, 102 #5 that are connected to each other via a network, and a host computer 101-2 having a remote copy manager RCM 105-2 ([0010]) and being connected with the secondary storage subsystems (Fig. 1), the method for remotely copying data from each of a plurality of primary volumes 110 to a corresponding secondary volume of a plurality of secondary volumes 117 of the invention (for example, the embodiment depicted in Fig. 1), as now recited in claim 1, involves primary volumes 110-1, 110-2, 110-3 that are constituted by the primary storage subsystems 102 #1, 102 #2, and the secondary volumes 117-1, 117-2, 117-3 that are constituted by the secondary storage subsystems 102 #3, 102 #4, 102 #5 (i.e., the consistency group C1 in Fig. 3). The method comprises (1) a normal synchronizing procedure (Figs. 4-5) and (2) an after-failure synchronizing procedure (Fig. 6-7).

The normal synchronizing procedure (Figs. 4-5) includes: receiving, at each of the secondary storage subsystems 102 #3, 102 #4, 102 #5, remote copy requests 120 in Fig. 1 (e.g., io-10 (T0), io-11 (T3), io-12 (T4), io-13 (T6), io-14 (T7) in 102 #3 in Fig. 5A) each of which is associated with a timestamp (e.g., T0, T3, T4, T6, T7) from each one of the plurality of primary storage subsystems 102 #1, 102 #2; receiving periodically (*“Each [primary] logical volume periodically issues a synchronize request 107 addressed to either all secondary logical volumes in the same consistency group or all secondary storage subsystems in the same consistency group.”*[0033]), at each of the secondary storage subsystems 102 #3, 102 #4, 102 #5, synchronizing requests 107 (e.g., SYNC-C1-#1-T7 in Fig. 5B) each of which is associated with a timestamp T7 and a primary storage ID #1 of a primary storage subsystem 102 #1, which sends a respective synchronizing request 107, from said each one of the primary storage subsystems 102 #1, 102 #2 respectively ([0063]); determining, at each of the secondary storage subsystems 102 #3, 102 #4, 102 #5, a first time as a first time parameter T5 based on the timestamp T7 included in the sync request 107; and determining, at each of the

secondary storage subsystems 102 #3, 102 #4, 102 #5, which remote copy requests 120 in Fig. 1 (e.g., io-10 (T0), io-11 (T3), io-12 (T4), io-13 (T6), io-14 (T7) in 102 #3) to process ([0063]) based on the first time parameter T5, primary storage IDs (e.g., “#1”) and timestamps (e.g., T0, T3, T4, T6, T7) associated with the remote copy requests, thereby maintaining data I/O consistency among said storage subsystems ([0007]).

The invention provides a unique remote-copy-request discriminative processing scheme in an N:M configuration ([0007]) of “*n primary storage subsystems and m secondary storage subsystems*” ([0004]) to maintain IO consistency in the consistency group.

For example, if the retrieved timestamp T7 is greater than the corresponding time parameter 111-1 #1 T3, an update occurs (As $T7 > T3$, the time parameter 111-1 #1 is updated from “T3” in the 102 #1 box in Fig. 5A to “T7” in the 102 #1 box in Fig. 5B). This update makes the time parameter 111-1 #1 T7 greater than the time parameter 111-1 #2 T5 in Fig. 5B, so that time parameter 112-1 is updated from “T3” in the 102 #1 box in Fig. 5A to “T5” in the 102 #1 box in Fig. 5B. When the time parameter 112-1 is updated, since the timestamp (T4) associated with io-12 is smaller than the timestamp (T5) that time parameter 112-1 indicates, only io-12 (T4) (*but not io-13 (T6), io-14 (T7)*), is moved to the disk request queue 116-1. As another example shown in Figs. 7A-C ([0074]), T3, T5 are determined to be the synchronized times for the secondary storage subsystems 102 #3-5. These changes cause the time parameters 112-2 and 112-3 to be updated to T3 such that io-23 (T2) and io-31 (T2) are moved to the disk request queues 116-2 and 116-3, respectively.

The after-failure synchronizing procedure (Figs. 6-7; “*FIG. 6 illustrates a flow diagram of synchronization procedure between the secondary storage subsystems after remote copy pair suspension*” [0018]; [0071]) includes: suspending said remote copy requests after a failure occurs (different types of failures are described in [0067]-[0070]); collecting and comparing by the RCM 105-2 in the host computer 101-2 (as shown in Fig. 1 a remote copy manager (RCM) 105-2 resides in the host computer 101-2; [0069]; “*when suspension occurs, RCM 105-2 collects the time indicated by the time parameters ...*” [0074]) time parameters 111 stored in the secondary storage systems to determine a synchronized time T3 (Steps 601-602 in Figs. 6A-B; [0074]); receiving from the RCM 105-2 in said host computer 101-2 at each of the secondary storage subsystems ([0071]) synchronizing requests 701 (Fig. 7B) each of which includes said synchronized time T3 (Steps 603-604); updating second time parameters 112 at each of the secondary storage subsystems up to said synchronized time T3 (Step 605; Fig. 7B); and determining, at each of the secondary storage subsystems, which remote copy requests to process based on said updated second time parameter 112 (Steps

605-608; Figs. 7C-G), thereby maintaining data I/O consistency among said storage subsystems.

For example, prior to the suspension, io-11 (T3) has already been moved to the disk request queue 116-1, but io-23 (T2) and io-31 (T2) are still in the suspended remote copy request queues 115-2 and 115-3, respectively. To achieve consistency among the secondary storage subsystems 102 #3-5, io-23 (T2) and io-31 (T2) are moved to the disk request queues 116-2 and 116-3, respectively ([0073]) as shown in Figs. 7A-G following the flow depicted in Fig. 6A or Fig. 6B. In FIG. 7A, among 1st time parameters 111-1 #1, 111-2 #1 and 111-3 #1, T3 is the greatest value such that T3 is determined by the RCM 105-2 of the host computer 101-2 to be the synchronized time for the secondary storage subsystem 102 #1 and included in a synchronization request 701 to be sent the secondary storage subsystem 102 #3-#5. This change causes the 2nd time parameters 112-2 and 112-3 to be updated to T3 in Fig. 7B. In FIG. 7C, any requests in remote copy request queues 115 with a time parameter smaller than or equal to T3, such as io-23 (T2) and io-31 (T2), are moved to the disk request queues 116-2 and 116-3, respectively (step 605 in FIG. 6A). What is left in the remote copy request queue 115 is then used to set the bitmap 114 in Fig. 7D to clear the remote copy request queue 115 as shown in Fig. 7E. Fig. 7E also shows queues in the queue 116 are performed to clear the queue 116 as shown in Fig. 7F. Fig. 7G shows a host accesses the secondary storage subsystems ([0073]-[0081]; claims 5-7, 12-14, 19-21).

The invention recited in claim 8 is directed to software for remotely copying data from a plurality of primary volumes to a plurality of secondary volumes according to the method recited in claim 1.

The invention recited in claim 15 is directed to a system for remotely copying data from each of a plurality of primary volumes to a plurality of secondary volumes according to the method recited in claim 1.

The invention keeps consistency of the stored data in a plurality of secondary storage systems based on the sync requests issued by the RCM in the host computer during a suspension period. (before the volume resynchronization between the primary and secondary systems) with the follow steps (Fig. 6A; [0074]):

- (1) the RCM in the host computer collects time indicated by the time parameters from the secondary storage systems when suspension between primary storage systems and secondary storage systems occurs (Step 601);

- (2) the RCM in the host computer compares the collected time so as to determine a synchronized time (Step 602);

(3) the RCM in the host computer issues sync request to the secondary storage systems (Step 603); and

(4) the secondary systems updates the secondary volumes based on the sync request from the host computer (Step 606).

Applicant respectfully contends that none of the cited prior art references teaches or suggests such an “after-failure synchronizing procedure executed by the RCM 105-2 in the host computer 101-2 and the secondary storage subsystems (after a failure occurs) in which remote copy requests being suspended” as in the present invention.

Crockett was relied upon by the Examiner (p. 8, line 5 to p. 10, line 8 of the outstanding Office Action) to teach the “after-failure synchronizing procedure” of the present invention. However, the alleged “after-failure synchronizing procedure” in Crockett (i.e., “volume resynchronization”) is executed by the primary SCU 27 (rather than “the RCM 105-2 residing in the host computer 101-2”) as admitted by the Examiner (p. 8, last line of the outstanding Office Action) and described on col. 10, lines 44-57 of Crockett (e.g., “*Referring now to FIGS. 8 and 9, there is shown the volume resynchronization process managed at the primary SCU 27.*” Col. 10, lines 58-60).

After suspending the primary volume 29 by signaling the counterpart SCUs (col. 10, lines 48-50), Crockett’s host CPU 2 simply resigns from the volume resynchronization process, and lets the primary and remote secondary storage subsystems (col. 10, lines 45-47) run the operations. Crockett’s general volume resynchronization process between the primary storage system and the secondary storage system starts with the primary storage subsystem according to Fig. 8 to read and send the groups of primary tracks modified during the suspension (col. 11, line 28) and their timestamp to the secondary SCU 31 (col. 10, line 44 to col. 11, line 29), and then the remote secondary storage subsystem writes out the groups of primary tracks modified before or during the suspension (col. 11, line 35) to the secondary volume in monotonic address and timestamp order according to Fig. 10 (col. 11, line 32-38).

Applicant contends that the cited references and their combinations all fail to teach or disclose each and every feature of the present invention as recited in independent claims 1, 8 and 15. As such, the present invention as now claimed is distinguishable and thereby allowable over the rejections raised in the Office Action. The withdrawal of the outstanding prior art rejections is in order, and is respectfully solicited.

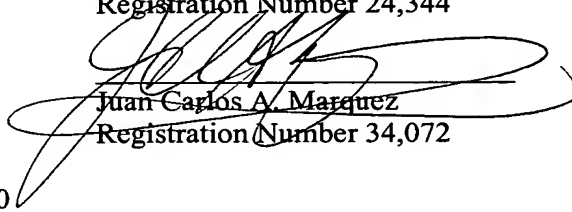
Conclusion

In view of all the above, Applicant respectfully submits that certain clear and distinct differences as discussed exist between the present invention as now claimed and the prior art references upon which the rejections in the Office Action rely. These differences are more than sufficient that the present invention as now claimed would not have been anticipated nor rendered obvious given the prior art. Rather, the present invention as a whole is distinguishable, and thereby allowable over the prior art.

Favorable reconsideration of this application as amended is respectfully solicited. Should there be any outstanding issues requiring discussion that would further the prosecution and allowance of the above-captioned application, the Examiner is invited to contact the Applicant's undersigned representative at the address and telephone number indicated below.

Respectfully submitted,

Stanley P. Fisher
Registration Number 24,344



Juan Carlos A. Marquez
Registration Number 34,072

REED SMITH LLP
3110 Fairview Park Drive, Suite 1400
Falls Church, Virginia 22042
(703) 641-4200

March 7, 2007

SPF/JCM/JT